

**CLAIMS:**

- 1           1.     A method for reporting topology changes in a subnet of a switched fabric  
2 including at least a client, a subnet manager (SM) and switches interconnected via links, said  
3 method comprising:  
4           creating and reporting a list of topology changes that are interesting to the client for  
5 topology change notifications;  
6           when a topology change occurs in the subnet, determining if the topology change is in the  
7 list of topology changes created by the interested client; and  
8           if the topology change is in the list of topology changes created by the interested client,  
9 reporting a topology change event to the interested client.
- 1           2.     The method as claimed in claim 1, wherein said list of topology changes is created  
2 by the client to serve as client-defined filters that specify the types of topology changes the client  
3 is interested in receiving notifications.
- 1           3.     The method as claimed in claim 2, wherein said list of topology changes includes,  
2 but is not limited to, when a new data path is created between a pair of end nodes in the subnet,  
3 when an existing data path is destroyed between a pair of end nodes in the subnet, when a new  
4 device is inserted in the subnet, and when an existing device is removed from the subnet.

1           4.       The method as claimed in claim 1, wherein said client corresponds to an end node  
2 of the subnet having at least one channel adapter (CA) installed to support one or more ports for  
3 data communication via said links of the subnet.

1           5.       The method as claimed in claim 2, wherein said determining the topology change  
2 in the list of topology changes and said reporting the topology change events to the interested  
3 client are executed by said subnet manager.

1           6.       The method as claimed in claim 5, wherein said subnet manager (SM) is installed  
2 in another end node of the subnet, and is implemented either in hardware or software to provide  
3 management services for all switches and end nodes in the subnet.

1           7.       The method as claimed in claim 5, wherein said subnet manager (SM) is installed  
2 in another end node of the subnet, and is implemented in software written using a high-level  
3 computer programming language for performing network management functions in compliance  
4 with the InfiniBand™ Architecture specification.

1           8.       The method as claimed in claim 5, wherein said subnet manager (SM) is installed  
2 in another end node of the subnet for discovering the subnet topology, assigning unique

1 addresses to all ports that are connected to the subnet, and establishing possible data paths among  
2 all ports by programing switch forwarding tables for download to the switches in the subnet for  
3 routing data packets to destinations via possible data paths established between switch pairs.

1 9. The method as claimed in claim 1, wherein said client sends a VendorSet  
2 (SetNotificationFilter) message to the subnet manager (SM) after the list of topology changes is  
3 created to indicate the topology changes that require client notifications, and said subnet manager  
4 (SM) sends a VendorGetResp (SetNotificationFilter) message back to the interested client to  
5 confirm receipt of the list of topology changes that the client is interested.

1 10. The method as claimed in claim 1, wherein said subnet manager (SM) sends a  
2 VendorSend (TopologyChangeNotification) message to the interested client after the topology  
3 change is determined in the list of topology changes to notify the topology change that occurred,  
4 and said client sends a VendorSendResp (TopologyChangeNotification) message back to the  
5 subnet manager (SM) to acknowledge the topology change notification.

1 11. A data network, comprising:  
2 a host system having at least one channel adapter (CA) installed therein supporting one or  
3 more ports;  
4 at least one target system having at least one channel adapter (CA) installed therein

1 supporting one or more ports;

2 a switched fabric comprising a plurality of different switches which interconnect said host  
3 system via CA ports to said remote system via CA port along different physical links for data  
4 communications; and

5 a fabric manager provided in said host system for making topology discovery, assigning  
6 local identifiers (LIDs) to all ports that are connected in the switched fabric, and programming  
7 forwarding tables for switches in the switched fabric, wherein said fabric manager includes a  
8 topology change notification mechanism configured to provide topology change notifications by:

9 enabling a client at one of the host system and the target system to create and  
10 communicate a list of topology changes that are interesting to the client for topology  
11 change notifications;

12 determining if a topology change occurred in the switched fabric is in the list of  
13 topology changes created by the interested client; and

14 if the topology change is in the list of topology changes created by the interested  
15 client, reporting a topology change event to the interested client.

1 12. The data network as claimed in claim 11, wherein said list of topology changes is  
2 created by the client to serve as client-defined filters that specify the types of topology changes  
3 the client is interested in receiving topology change notifications.

1           13.     The data network as claimed in claim 12, wherein said list of topology changes  
2 includes, but is not limited to, when a new data path is created between a pair of end nodes in the  
3 switched fabric, when an existing data path is destroyed between a pair of end nodes in the  
4 switched fabric, when a new device is inserted in the switched fabric, and when an existing  
5 device is removed from the switched fabric.

1           14.     The data network as claimed in claim 11, wherein said fabric manager is installed  
2 in another one of the host system and the target system, and is implemented either in hardware or  
3 software to provide management services for all switches and end nodes in the switched fabric.

1           15.     The data network as claimed in claim 11, wherein said fabric manager is installed  
2 in another one of the host system and the target system, and is implemented in software written  
3 using a high-level computer programming language for performing network management  
4 functions in compliance with the InfiniBand™ Architecture specification.

1           16.     The data network as claimed in claim 15, wherein said fabric manager is further  
2 configured to discover the fabric topology, assign unique addresses to all ports that are connected  
3 to the switched fabric, and establish possible data paths among all ports by programing switch  
4 forwarding tables for download to the switches in the switched fabric for routing data packets to  
5 destinations via possible data paths established between switch pairs.

1           17.    The data network as claimed in claim 11, wherein said client sends a VendorSet  
2    (SetNotificationFilter) message to the fabric manager after the list of topology changes is created  
3    to indicate the topology changes that require client notifications, and said fabric manager sends a  
4    VendorGetResp (SetNotificationFilter) message back to the interested client to confirm receipt of  
5    the list of topology changes that the client is interested.

1           18.    The data network as claimed in claim 11, wherein said fabric manager sends a  
2    VendorSend (TopologyChangeNotification) message to the interested client after the topology  
3    change is determined in the list of topology changes to notify the topology change that occurred,  
4    and said client sends a VendorSendResp (TopologyChangeNotification) message back to the  
5    fabric manager to acknowledge the topology change notification.

1           19.    A computer readable medium comprising instructions that, when executed by a  
2    host system in a switched fabric including end nodes and switches interconnected via links, cause  
3    the host system to:

4           enabling a client at an end node to create and communicate a list of topology changes that  
5    are interesting to the client for topology change notifications;

6           determining if a topology change occurred in the switched fabric is in the list of topology  
7    changes created by the interested client; and

1 if the topology change is in the list of topology changes created by the interested client,  
2 reporting a topology change event to the interested client.

1 20. The computer readable medium as claimed in claim 19, wherein said list of  
2 topology changes is created by the client to serve as client-defined filters that specify the types of  
3 topology changes the client is interested in receiving topology change notifications.

1 21. The computer readable medium as claimed in claim 20, wherein said list of  
2 topology changes includes, but is not limited to, when a new data path is created between a pair  
3 of end nodes in the switched fabric, when an existing data path is destroyed between a pair of end  
4 nodes in the switched fabric, when a new device is inserted in the switched fabric, and when an  
5 existing device is removed from the switched fabric.

1 22. The computer readable medium as claimed in claim 19, further causing the system  
2 to enable the client to send a VendorGetResp (SetNotificationFilter) message to the interested  
3 client upon receipt of a VendorSet (SetNotificationFilter) message from the interested client to  
4 confirm receipt of the list of topology changes that the client is interested.

1 23. The computer readable medium as claimed in claim 19, further causing the system  
2 to send a VendorSend (TopologyChangeNotification) message to the interested client after the

- 1 topology change is determined in the list of topology changes to notify the topology change that
- 2 occurred, and to acknowledge the topology change notification upon receipt of a
- 3 VendorSendResp (TopologyChangeNotification) message from the interested client.